

Preserving multiple first integrals by discrete gradients

Morten Dahlby Brynjulf Owren Takaharu Yaguchi

June 9, 2011

Abstract

We consider systems of ordinary differential equations with known first integrals. The notion of a discrete tangent space is introduced as the orthogonal complement of an arbitrary set of discrete gradients. Integrators which exactly conserve all the first integrals simultaneously are then defined. In both cases we start from an arbitrary method of a prescribed order (say, a Runge-Kutta scheme) and modify it using two approaches: one based on projection and one based on local coordinates. The methods are tested on the Kepler problem.

1 Introduction

A system of ordinary differential equations which preserves a first integral $H(y)$ can be written in the form

$$\dot{y} = f(y) = S(y)\nabla H(y), \quad y \in \mathbb{R}^m, \quad (1.1)$$

where $S(y)$ is an antisymmetric matrix, see [12]. An approximate numerical solution, $y^n \approx y(t^n)$, $n \geq 1$, is said to be integral preserving if $H(y^n) = H(y^0)$, $n \geq 1$. Energy preserving methods go all the way back to the seminal paper by Courant, Friedrichs, Lewy [4], where an energy preserving difference scheme is derived and the property is used to prove convergence of the scheme. Recently, it has become increasingly popular to study conservative schemes as discrete dynamical systems in their own right, attempting to mimic properties of the continuous system by introducing suitable discrete counterparts. An example of importance in this note is the replacement of the gradient in (1.1) by a discrete gradient operator.

The discrete gradient method is one of the most prevalent approaches in the literature. It was first systematically treated by Gonzalez [5] and McLachlan et al. [12]. The idea is to introduce a discrete approximation to the gradient, letting $\bar{\nabla}H : \mathbb{R}^m \times \mathbb{R}^m \rightarrow \mathbb{R}^m$ be a continuous map satisfying

$$\begin{aligned} H(u) - H(v) &= \bar{\nabla}H(v, u)^\top (u - v), \\ \bar{\nabla}H(u, u) &= \nabla H(u). \end{aligned}$$

The existence of such discrete gradients is well established in the literature, see for instance the monograph by Hairer et al. [6] or the papers [5] and [12]. Their

construction is not unique, we give here two different examples. The Averaged Vector Field (AVF) gradient [17] is defined as

$$\bar{\nabla}_{\text{AVF}} H(v, u) = \int_0^1 \nabla H(\xi u + (1 - \xi)v) d\xi. \quad (1.2)$$

The coordinate increment method [7] is defined in terms of the coordinates of the vectors v and u , the i th component of $\bar{\nabla} H(v, u)$ is then given as

$$(\bar{\nabla}_{\text{CI}} H(v, u))_i = \frac{H(u_1, \dots, u_i, v_{i+1}, \dots, v_m) - H(u_1, \dots, u_{i-1}, v_i, \dots, v_m)}{u_i - v_i}. \quad (1.3)$$

An important difference between these two discrete gradients is that (1.2) is symmetric, $\bar{\nabla}_{\text{AVF}} H(v, u) = \bar{\nabla}_{\text{AVF}} H(u, v)$, while (1.3) is not. However, note that a symmetric version of the coordinate increment discrete gradient can be constructed by

$$\bar{\nabla}_{\text{SCI}} H(v, u) = \frac{1}{2} (\bar{\nabla}_{\text{CI}} H(v, u) + \bar{\nabla}_{\text{CI}} H(u, v)). \quad (1.4)$$

Once a discrete gradient has been found, one immediately obtains an integral preserving method by simply letting

$$\frac{y^{n+1} - y^n}{h} = \bar{S}(y^n, y^{n+1}) \bar{\nabla} H(y^n, y^{n+1}).$$

Here h is the time step, and $\bar{S}(y^n, y^{n+1})$ is some skew-symmetric approximation to the matrix S , one would normally require that $S(y) = \bar{S}(y, y)$. We remark that discrete gradient methods are implicit.

In this note we consider the case where there are more than one first integral and the objective is to preserve any number of such invariants simultaneously. Some earlier attempts to achieve this include the papers [12, 18] in which the antisymmetric matrix $S(y)$ is replaced by an antisymmetric tensor taking discrete gradients of all integrals to be preserved as input. A formula for this antisymmetric tensor is given. Another approach is an integrator for a class of separable Hamiltonian systems ([14] and references therein), where the integrator which preserves all integrals is designed based on separation of variables by the Kustaanheimo–Stiefel transformation [20]. This transformation was also adopted in the development of the “exact” integrator for the Kepler problem by Kozlov ([8], see also [2]). It would be also noteworthy that Labudde and Greenspan [9] proposed an energy-and-angular-momentum-preserving integrator for the differential equations of motion of classical mechanics, and they developed similar integrators of high order of convergence in the sequels [10, 11]. Energy and linear momentum are exactly conserved, but for the system case [11] angular momentum is not. Another approach is used by Simo et al. in [19] to develop schemes that preserve energy and momentum.

We shall instead present an approach which does not rely on finding such a tensor nor a structure of the equation, we only assume knowledge of the first integrals to be preserved as well as the ODE vector field itself. Examples of such invariants are the energy and momentum, but our approach is not limited to these. The discrete gradients are essential to the algorithm we develop, but note that the methods themselves are not discrete gradient methods in the

usual sense. After defining the general method, we present two particular cases based on projection and local coordinates, respectively. Both use an underlying scheme of arbitrary order p , we prove that the resulting schemes retain this order, regardless of the choice of discrete gradient.

In Section 3 we apply the new methods to the Kepler problem, a system with four degrees of freedom and three independent first integrals. We illustrate our approach by preserving combinations of one or more of these three integrals. An interesting question is whether the present methods based on discrete gradients perform better than the more standard projection methods which make use of the exact gradients of the first integrals. We show two examples where the different approaches are compared. In the examples we use Runge-Kutta methods of different order as the underlying schemes. The Kepler problem is of course a well-known and popular test case, and methods which preserve one or more integrals for this particular problem can be found in e.g. [2, 3, 8, 13], the methods in these references are derived by means of the Kustaanheimo–Stiefel transformation.

2 Preserving multiple invariants

Suppose that an ODE system (1.1) possesses $q \geq 1$ independent first integrals, $H_1(y), \dots, H_q(y)$. These invariants foliate \mathbb{R}^m into $(m - q)$ -dimensional submanifolds (leaves)

$$M = M_c = \{y \in \mathbb{R}^m : H_1(y) = c_1, H_2(y) = c_2, \dots, H_q(y) = c_q\}.$$

The tangent space $T_y M$ of M at y is the orthogonal complement to

$$\text{span}\{\nabla H_1(y), \dots, \nabla H_q(y)\}.$$

For simplicity we write only M for M_c for the rest of this paper.

Definition 2.1. Let $\bar{\nabla}$ be a fixed discrete gradient operator and let H_1, \dots, H_q be independent first integrals. The discrete tangent space at $(u, v) \in \mathbb{R}^m \times \mathbb{R}^m$ is

$$T_{(v,u)} M = \{\eta \in \mathbb{R}^m : \langle \bar{\nabla} H_j(v, u), \eta \rangle = 0, \ 1 \leq j \leq q\}.$$

A vector $\eta = \eta_{(v,u)} \in T_{(v,u)} M$ is called a discrete tangent vector.

Note that this definition causes $T_{(y,y)} M = T_y M$.

Lemma 2.2. Any integrator satisfying

$$y^{n+1} - y^n = \eta_{(y^n, y^{n+1})} \in T_{(y^n, y^{n+1})} M$$

preserves all integrals, in the sense that $H_i(y^{n+1}) = H_i(y^n)$, $1 \leq i \leq q$.

Proof. For any i we compute

$$\begin{aligned} H_i(y^{n+1}) - H_i(y^n) &= \bar{\nabla} H_i(y^n, y^{n+1})^\top (y^{n+1} - y^n) \\ &= \bar{\nabla} H_i(y^n, y^{n+1})^\top \eta_{(y^n, y^{n+1})} = 0. \end{aligned}$$

□

This way of devising integral preserving schemes is similar, but slightly different to the presentation in e.g. [6], [12] and [17]. In this paper we will outline two ways of ensuring that the condition of Lemma 2.2 is satisfied – projection and local coordinates. We emphasise, however, that there are other ways of satisfying Lemma 2.2. One example taken from [12] is

$$\eta_{(y^n, y^{n+1})} = h\bar{S}(y^n, y^{n+1})\bar{\nabla}H_1(y^n, y^{n+1})\dots\bar{\nabla}H_q(y^n, y^{n+1}),$$

where \bar{S} is a q -dimensional skew-symmetric tensor.

2.1 Projection

We consider (1.1) with the first integrals $H_1(y), \dots, H_q(y)$. We propose the projection scheme

$$u^{n+1} = \phi_h(y^n), \quad y^{n+1} = y^n + \mathcal{P}(y^n, y^{n+1})(u^{n+1} - y^n) \quad (2.1)$$

where ϕ_h is the discrete flow that defines an arbitrary method of order p

$$y(t+h) - u^{n+1} = y(t+h) - \phi_h(y(t)) = \mathcal{O}(h^{p+1}),$$

and $\mathcal{P}(y^n, y^{n+1})$ is a smooth projection operator onto the discrete tangent space $T_{(y^n, y^{n+1})}M$. An alternative method is

$$y^{n+1} = y^n + h\mathcal{P}(y^n, y^{n+1})\psi_h(y^n, y^{n+1}) \quad (2.2)$$

where ψ_h is the increment function that defines a method of the form

$$y^{n+1} = y^n + h\psi_h(y^n, y^{n+1}).$$

This method is itself assumed to be of order p , that is

$$y(t+h) - y(t) - h\psi_h(y(t), y(t+h)) = \mathcal{O}(h^{p+1}). \quad (2.3)$$

We remark in passing that if both the discrete gradients $\bar{\nabla}H_i$ and the increment function ψ_h are symmetric, then the method (2.2) is symmetric.

Example 2.3. *Using Runge-Kutta as the underlying scheme ϕ_h we can construct examples of (2.1). The unprojected solution is given as*

$$u^{n+1} = y^n + h \sum_{i=1}^s b_i k_i,$$

where k_1, \dots, k_s are the solutions to the (possibly implicit) equations

$$k_i = f \left(y^n + h \sum_{j=1}^s a_{ij} k_j \right).$$

Even if the Runge-Kutta scheme is explicit the scheme (2.1) will be implicit since y^{n+1} appears in the projection operator

$$y^{n+1} = y^n + h\mathcal{P}(y^n, y^{n+1}) \left(\sum_{i=1}^s b_i k_i \right).$$

The difference between (2.1) and (2.2) is subtle, but to illustrate that they are in fact distinct we consider the implicit midpoint method as the underlying scheme and we get for the two methods

$$\begin{aligned} y^{n+1} &= y^n + h\mathcal{P}(y^n, y^{n+1})f\left(\frac{y^n + u^{n+1}}{2}\right), \\ y^{n+1} &= y^n + h\mathcal{P}(y^n, y^{n+1})f\left(\frac{y^n + y^{n+1}}{2}\right), \end{aligned}$$

where in the former method u^{n+1} is computed by solving

$$u^{n+1} = y^n + hf\left(\frac{y^n + u^{n+1}}{2}\right).$$

Theorem 2.4. *The schemes (2.1) and (2.2) are of order p , that is*

$$\begin{aligned} y(t+h) - y(t) - \mathcal{P}(y(t), y(t+h))(u^{n+1} - y(t)) &= \mathcal{O}(h^{p+1}), \\ u^{n+1} &= \phi_h(y(t)), \end{aligned} \quad (2.4)$$

and

$$y(t+h) - y(t) - h\mathcal{P}(y(t), y(t+h))\psi_h(y(t), y(t+h)) = \mathcal{O}(h^{p+1}). \quad (2.5)$$

Proof. We use the shorthand notation \mathcal{P} for $\mathcal{P}(y(t), y(t+h))$ in this proof. To prove (2.4), we compute

$$\begin{aligned} &y(t+h) - y(t) - \mathcal{P}(u^{n+1} - y(t)) \\ &= y(t+h) - y(t) - (\mathcal{I} - (\mathcal{I} - \mathcal{P}))(u^{n+1} - y(t)) \\ &= y(t+h) - y(t) - (u^{n+1} - y(t)) + (\mathcal{I} - \mathcal{P})(u^{n+1} - y(t)). \end{aligned}$$

Since ϕ_h is of order p , we have

$$y(t+h) - y(t) - (u^{n+1} - y(t)) = y(t+h) - u^{n+1} = \mathcal{O}(h^{p+1}). \quad (2.6)$$

Therefore if we have

$$(\mathcal{I} - \mathcal{P})(u^{n+1} - y(t)) = \mathcal{O}(h^{p+1}),$$

the proof is completed. This estimate is obtained in the following way. Because the image of $\mathcal{I} - \mathcal{P}(y(t), y(t+h))$ is spanned by

$$\{\bar{\nabla}H_1(y(t), y(t+h)), \dots, \bar{\nabla}H_q(y(t), y(t+h))\},$$

it is enough to show

$$\bar{\nabla}H_i(y(t), y(t+h)) \cdot (u^{n+1} - y(t)) = \mathcal{O}(h^{p+1}).$$

From (2.6) we obtain

$$u^{n+1} - y(t) = y(t+h) - y(t) + \mathcal{O}(h^{p+1}),$$

and hence

$$\begin{aligned}
& \bar{\nabla} H_i(y(t), y(t+h)) \cdot (u^{n+1} - y(t)) \\
&= \bar{\nabla} H_i(y(t), y(t+h)) \cdot (y(t+h) - y(t) + \mathcal{O}(h^{p+1})) \\
&= H_i(y(t+h)) - H_i(y(t)) + \mathcal{O}(h^{p+1}) \\
&= \mathcal{O}(h^{p+1}).
\end{aligned}$$

The last equality is from the conservation property of the original equation. The proof of (2.5) is almost identical and therefore omitted. \square

We remark that in what we call standard projection methods u^{n+1} is projected orthogonally onto the manifold M by computing

$$\min \|y^{n+1} - u^{n+1}\| \quad \text{subject to } y^{n+1} \in M.$$

This can be achieved for instance by using Lagrange multipliers (see [6], for instance). This approach differs from ours since we project along the discrete gradients which depend on the end point y^{n+1} .

Computing the projector. A simple and straightforward way of obtaining the projector $\mathcal{P}(y^n, y^{n+1})$ is as follows: Define the $(q \times m)$ -matrix $Y = Y(y^n, y^{n+1})$ whose columns are the discrete gradients $\bar{\nabla} H_i(y^n, y^{n+1})$, $i = 1, \dots, q$. Compute a reduced QR -decomposition $QR = Y$ where $Q \in \mathbb{R}^{m \times q}$ and $R \in \mathbb{R}^{q \times q}$. Then define the projection matrix as $\mathcal{P}(y^n, y^{n+1}) = I - QQ^\top$.

2.2 Local coordinates

The local coordinates approach presented here is basically of the same type as the standard method by Potra and Rheinboldt ([15], see also [6, 16]). One important difference is that our local coordinates are algorithmically constructed by using discrete gradients. We also present an "automatic differentiation" algorithm of the coordinate map, which could be used to increase the efficiency of the computations.

Inspired by [1], we consider local coordinates on a chart containing y^0 by defining a map $\eta \mapsto y = \chi(\eta)$. The map is defined implicitly by

$$\chi(\eta) = y : y - y^0 = T(y^0, y)\eta, \quad (2.7)$$

where $T(y^0, y)$ is a smooth $m \times (m - q)$ -matrix whose columns form a basis for the left nullspace (orthogonal column complement) of the matrix $Y(y^0, y) = [\bar{\nabla} H_1(y^0, y), \dots, \bar{\nabla} H_q(y^0, y)]$. We suppress the dependency on y^0 and use the shorthand notation $T(y)$ and $Y(y)$ for the rest of this paper.

Lemma 2.5. *Suppose that $\nabla H_1(y^0), \dots, \nabla H_q(y^0)$ are linearly independent for all $y^0 \in M$. Suppose also that for all $y^0 \in M$, $\bar{\nabla} H_1(y^0, y), \dots, \bar{\nabla} H_q(y^0, y)$ are C^∞ with respect to y . Then the following statements hold.*

1. (2.7) defines a one-to-one map ω_{y^0} in a neighborhood $N_{y^0} \subset M$. ω_{y^0} and $\omega_{y^0}^{-1}$ are C^∞ .
2. The collection of the pairs $\{(N_{y^0}, \omega_{y^0}) \mid y^0 \in M\}$ forms an atlas of M .

Proof. 1. From the continuity of the discrete gradients, we deduce that for all $y^0 \in M$, there exists a neighborhood \tilde{N}_{y^0} in \mathbb{R}^m of y^0 in which $\bar{\nabla}H_1(y^0, y), \dots, \bar{\nabla}H_q(y^0, y)$ are linearly independent. For $y \in \tilde{N}_{y^0}$, $T(y)$ admits the QR decomposition $T(y) = QR$ and η is obtained by $\eta = (R^\top R)^{-1}Q^\top(y - y^0) = (T^\top T)^{-1}T^\top(y - y^0)$. This is a C^∞ function. Conversely, the Jacobian matrix of the function $\eta(y)$ at $y = y^0$ is

$$\frac{\partial \eta}{\partial y}(y^0) = (T^\top T)^{-1}T^\top$$

and hence

$$\text{rank} \frac{\partial \eta}{\partial y}(y^0) = \text{rank}(T^\top T)^{-1}T^\top = m - q.$$

Thus ω is defined in a neighborhood \bar{N}_{y^0} of y^0 by the implicit function theorem and is C^∞ . The proof is completed by letting $N_{y^0} = \bar{N}_{y^0} \cap \tilde{N}_{y^0} \cap M$.

2. This is immediately obtained from the first statement. \square

Consider now the curve $\eta(t)$ and let $y(t) = \chi(\eta(t))$. We differentiate the curve to obtain from (2.7)

$$\dot{y}(t) = T'_{y(t)}(\dot{\eta}(t))\eta(t) + T(y(t))\dot{\eta}(t).$$

From this we compute

$$\dot{\eta} = -T^\top(\chi \circ \eta)T'_{\chi \circ \eta}(f(\chi \circ \eta))\eta + T^\top(\chi \circ \eta)f(\chi \circ \eta) \quad (2.8)$$

where the original ODE is $\dot{y} = f(y)$. The method we propose takes one step as follows

1. Let $\eta_0 = 0$.
2. Take a step with any p th order method applied the ODE (2.8) using $y^0 = y^n$ in (2.7). The result is η_1 .
3. Compute $y^{n+1} = \chi(\eta_1)$.

We immediately obtain the next theorem from Lemma 2.5, because the solution curve lies in M and a p th order method is applied in a chart of M .

Theorem 2.6. *Under the assumptions of Lemma 2.5, the above scheme is of order p .*

The main difficulty in this approach is the computation of the derivative map $T'_y(\zeta)$ for arbitrary values of $y \in \mathbb{R}^m$ and $\zeta \in \mathbb{R}^m$. This is needed explicitly in the integration algorithm, but may also be a useful tool in computing the coordinate map (2.7). We may define $T(y)$ as the last $m - q$ columns of the $m \times m$ -matrix $Q(y)$ defined through a QR-decomposition where $Y(y) = Q(y)R(y)$ and where we have used the shorthand notation

$$Y(y) = [\bar{\nabla}H_1(y^n, y), \dots, \bar{\nabla}H_q(y^n, y)]$$

We realise the QR-decomposition by means of the Householder method, applying a sequence of q elementary orthogonal transformations to the matrix $Y(y)$ as described in most elementary text books in numerical linear algebra, see e.g [21]. Each transformation is of the form

$$Q_k = I - 2v_k v_k^\top, \quad v_k \in \mathbb{R}^m, \quad v_k^\top v_k = 1,$$

and its aim is to eliminate all elements under the diagonal in the k th column of the matrix to which it is applied.

In order to explain how we compute the derivative $Q'_y(\zeta) =: DQ(y, \zeta)$, we first review the Householder method.

```

 $Y^{(1)} := Y$ 
for  $k = 1 : q$ ,
   $w_k = \Pi_k Y_k^{(k)} - \|\Pi_k Y_k^{(k)}\| e_k$ 
   $v_k = \frac{w_k}{\|w_k\|}$ 
  for  $r = k : q$ ,
     $Y_r^{(k+1)} = (I - 2 v_k v_k^\top) Y_r^{(k)}$ 
  end
end

```

where the following conventions have been used:

- $\|\cdot\|$ is the Euclidean norm.
- $Y^{(k+1)} = Q_k Y^{(k)}$, $Y_r^{(k)}$ is column r of $Y^{(k)}$.
- The projector Π_k puts zeros in the first $k - 1$ components and leaves the rest of the components unchanged when applied to a vector in \mathbb{R}^m .
- e_k is the k th canonical unit vector in \mathbb{R}^m .

The vectors v_k computed in the algorithm contain all information needed to reconstruct the factor Q , whereas $R := Y^{(q+1)}$. For simplicity, and to avoid the loss of regularity in $Q(y)$ viewed as a matrix valued function of y , we have here ignored the sign convention which is usually applied in the definition of w_k [21]. The idea is now to differentiate the variables in the algorithm with respect to y , writing for any object, say $X(y)$, its derivative as $DX = DX(y, \zeta) = \frac{d}{d\varepsilon} \Big|_{\varepsilon=0} X(y + \varepsilon \zeta)$ for any $y, \zeta \in \mathbb{R}^m$. The dependence on y, ζ will usually be suppressed in the notation when no confusion is at risk. Notice that D commutes with Π_k for any k . The following recursion formulae are easily derived

$$\begin{aligned}
Dw_k &= \Pi_k DY_k^{(k)} - \frac{(\Pi_k Y_k^{(k)})^\top \Pi_k DY_k^{(k)}}{\|\Pi_k Y_k^{(k)}\|} e_k, \\
Dv_k &= \left(Dw_k - \frac{w_k^\top Dw_k}{\|w_k\|^2} w_k \right) \|w_k\|^{-1}, \\
DY_r^{(k+1)} &= DY_r^{(k)} - 2 \left(v_k^\top Y_r^{(k)} Dv_k + Dv_k^\top Y_r^{(k)} v_k + v_k^\top DY_r^{(k)} v_k \right).
\end{aligned}$$

The initial $DY^{(1)} = DY^{(1)}(y, \zeta)$ should be computed by differentiating the q discrete gradients $\bar{\nabla} H_1(p, y), \dots, \bar{\nabla} H_q(p, y)$ with respect to y . When the discrete

gradients are given by the AVF formula, we may derive the following expressions for the r th column of $DY^{(1)}$

$$DY_r^{(1)}(y, \zeta) = D\bar{\nabla}H_r(y, \zeta) = \left(\int_0^1 \xi \nabla^2 H_r(\xi y + (1 - \xi)p) d\xi \right) \cdot \zeta$$

where $\nabla^2 H_r$ is the Hessian of the integral H_r . The expression for the coordinate increment cases are given in the appendix.

In the present application we only make use of the Q -part of the QR-decomposition, thus we need only store v_1, \dots, v_q and Dv_1, \dots, Dv_q for subsequent use. We may summarize the extended algorithm for computing these quantities as follows, using a Matlab inspired indexing notation where the sub-matrix $Y_{a:b,c:d}$ of Y means

$$Y_{a:b,c:d} = \begin{pmatrix} Y_{a,c} & \cdots & Y_{a,d} \\ \vdots & \ddots & \vdots \\ Y_{b,c} & \cdots & Y_{b,d} \end{pmatrix}.$$

Given Y and DY as $m \times q$ -matrices

for $k = 1 : q$

$$w = Y_{k:m,k} - \|Y_{k:m,k}\| e_1$$

$$Dw = DY_{k:m,k} - \frac{Y_{k:m,k}^\top DY_{k:m,k}}{\|Y_{k:m,k}\|} e_1$$

$$\tilde{v}_k = w / \|w\|$$

$$D\tilde{v}_k = (Dw - \frac{w^\top Dw}{\|w\|^2} w) \|w\|^{-1}$$

$$DY_{k:m,k+1:q} = DY_{k:m,k+1:q} - 2\tilde{v}_k \tilde{v}_k^\top DY_{k:m,k+1:q} - 2D\tilde{v}_k \tilde{v}_k^\top Y_{k:m,k+1:q} \\ - 2\tilde{v}_k D\tilde{v}_k^\top Y_{k:m,k+1:q}$$

$$Y_{k:m,k+1:q} = Y_{k:m,k+1:q} - 2\tilde{v}_k \tilde{v}_k^\top Y_{k:m,k+1:q}$$

end

The first $k - 1$ entries of the vectors v_k , Dv_k are zeros, and the remaining $m - k + 1$ entries are contained in \tilde{v}_k , $D\tilde{v}_k$ on exit. The complexity of this algorithm is $\mathcal{O}(mq^2 + q^3)$.

One should also note that we always multiply Q (DQ resp) by vectors $\bar{\eta} \in \mathbb{R}^m$ whose first q columns are zero, this may be taken advantage of in the implementation. The procedure for computing $Q\bar{\eta}$ by means of v_1, \dots, v_k is described in [21, Algorithm 10.3], the cost is $\mathcal{O}(mq)$. We may extend this algorithm so that it computes also $DQ\bar{\omega}$ given Dv_1, \dots, Dv_q . Defining the $m \times m$ -matrix $P_k = Q_k \cdots Q_q$, $k = 1, \dots, q$, we get the downwards recursion

$$P_{k-1} = Q_{k-1}P_k, \quad P_q = Q_q, \quad P_1 = Q = Q_1 \cdots Q_q,$$

and differentiation yields

$$DP_{k-1} = DQ_{k-1}P_k + Q_{k-1}DP_k \\ = -2(Dv_{k-1} v_{k-1}^\top + v_{k-1} Dv_{k-1}^\top)P_k + (I - 2v_{k-1} v_{k-1}^\top)DP_k$$

The following algorithm results for computing $\phi = DQ\bar{\omega}$

```

 $\psi = \bar{\omega}, \phi = 0$ 
for  $k = q : -1 : 1$ 
     $\phi = \phi - 2(v_k^\top \psi Dv_k + Dv_k^\top \psi v_k + v_k^\top \phi v_k)$ 
     $\psi = \psi - 2v_k^\top \psi v_k$ 
end

```

The complexity of this algorithm is $\mathcal{O}(mq)$.

3 Numerical integration of the Kepler problem

The Kepler two-body problem describes the motion of two bodies which attract each other. By placing the first body in the origin, the position (y_1, y_2) and the velocity (y_3, y_4) of the other body are given by the following four-dimensional ODE

$$\begin{aligned}
 \dot{y}_1 &= y_3, \\
 \dot{y}_2 &= y_4, \\
 \dot{y}_3 &= -\frac{y_1}{(y_1^2 + y_2^2)^{3/2}}, \\
 \dot{y}_4 &= -\frac{y_2}{(y_1^2 + y_2^2)^{3/2}}.
 \end{aligned} \tag{3.1}$$

This system preserves the Hamiltonian

$$H_1(y) = \frac{1}{2} (y_3^2 + y_4^2) - \frac{1}{\sqrt{y_1^2 + y_2^2}},$$

the angular momentum

$$H_2(y) = y_1 y_4 - y_2 y_3,$$

and the Runge-Lenz-Pauli vector

$$\begin{aligned}
 H_3(y) &= y_2 y_3^2 - y_1 y_3 y_4 - \frac{y_2}{\sqrt{y_1^2 + y_2^2}}, \\
 H_4(y) &= y_1 y_4^2 - y_2 y_3 y_4 - \frac{y_1}{\sqrt{y_1^2 + y_2^2}}.
 \end{aligned}$$

Since $q = m = 4$ any subset of three out of the four invariants is dependent. We want to compare schemes that preserve none, one, two, and all of the invariants above. We use the projection method (2.1) with the standard fourth order explicit Runge-Kutta method as the underlying scheme. The discrete gradients are calculated using (1.4). The schemes that preserve one of H_1, H_3 are denoted as RK4Proj1 and RK4Proj3, respectively. The scheme that preserves both H_1 and H_3 is called RK4Proj13. The original Runge-Kutta scheme preserves neither and is denoted RK4. The RK123Proj scheme is omitted from the plot since it produces exactly the ellipsis of the Kepler problem.

The resulting plots from these methods in Figure 3.1 are arranged according to the table

RK4	RK4Proj1
RK4Proj3	RK4Proj13

Table 3.1: The location of the schemes in the plots of Figure 3.1.

The initial values are taken from section I.2.3 of [6],

$$y_1^0 = 1 - e, \quad y_2^0 = 0, \quad y_3^0 = 0, \quad y_4^0 = \sqrt{\frac{1+e}{1-e}},$$

where the eccentricity is $e = 0.6$ and the exact solution has period 2π . The time step is $h = 0.2$ and we integrate for 50000 steps.

Figure 3.1 shows the numerical solutions. RK4 spirals inwards until it eventually blows up. RK4Proj1 has a counterclockwise precession effect. The Runge-Lenz-Pauli vector has to do with the orientation of the ellipse and RK4Proj3 does therefore not exhibit this effect, it will however converge to a smaller circle around the origin. The solution of RK4Proj13 shows an improvement compared to RK4Proj1 and RK4Proj3.

This example illustrates that there are cases where the preservation of one or more invariants are important to get a numerical solution with good long term properties. Not surprisingly, one observes a gradual improvement in the quality of the solution as the number of preserved first integrals increases. The extra computational effort needed to preserve multiple integrals compared to one is almost negligible, in this example the computation took less than 10% longer.

In Figure 3.2 we plot the global error of four schemes (RK2Proj123, RK4Proj123, RK5Proj123, and RK7Proj123) that preserve the four invariants. The underlying schemes are four RK-schemes of order 2, 4, 5, and 7. We see that the schemes attain the order of the underlying scheme, which is what we proved in Theorem 2.4.

Figure 3.3 shows a comparison between a standard orthogonal projection method as defined at the end of Section 2.1 and the projection method (2.1). Both preserve H_1 and H_2 simultaneously and use the implicit midpoint method as the underlying scheme. Figure 3.3a shows that our projection method is more accurate for $e = 0.6$ while Figure 3.3b shows that the standard projection method is more accurate for $e = 0.7$. In the standard projection method one has to compute the distance to the underlying manifold, usually denoted by the Lagrange multiplier λ , however for our projection scheme this is already (implicitly) known. The resulting nonlinear system will have dimension $m + q$ (see section IV.4 in [6]) compared to m for our proposed method. Our implementation uses the same nonlinear solver (Matlab's fsolve) for both methods. In Figure 3.3c we have adjusted h such that both schemes have the same computation time, in which case one sees that the new method is slightly better than the standard projection method even for $e = 0.7$.

Conclusion and further work. We have presented a new methodology for preserving multiple first integrals in systems of ordinary differential equations, using discrete gradients as the underlying tool. By using the notion of a discrete tangent space, two methodologies for designing numerical schemes are easily derived, projection and local coordinates. The resulting algorithms are relatively inexpensive compared to well-known algorithms preserving precisely one first




Figure 3.1: The numerical solution (thin line) of the Kepler problem (3.1) using the schemes of Table 3.1 with $h = 0.2$. The first 500 steps are shown. The exact solution (thick line) is an ellipse with eccentricity $e = 0.6$.







Figure 3.2: The global error of the four schemes RK2Proj123, RK4Proj123, RK5Proj123, and RK7Proj123. The dotted lines are reference lines of exact order.



(a) $e = 0.6$ and $h = 0.1$.



(b) $e = 0.7$ and $h = 0.05$.



(c) $e = 0.7$ and $h = 0.075$ (left) $h = 0.05$ (right).

Figure 3.3: The numerical solution (thin line) of the Kepler problem (3.1) using the midpoint method with standard orthogonal projection (left) and (2.1) with the midpoint method as the underlying scheme. The exact solution (thick line) is an ellipse with eccentricity e . Both schemes preserves H_1 and H_2 exactly. In the last row (c) the step size h is adjusted such that both schemes have the same computation time.

integral. There are of course several other well known methods for preserving multiple invariants, but we believe that the new methods which are based on discrete gradients rather than exact ones may be attractive for certain problems. Symmetric schemes are easily constructed, no Lagrange multipliers are required, and the schemes can be seen as natural generalizations of the popular discrete gradient methods. Although the present paper considers only systems of ordinary differential equations, the approach taken may be easily adapted to partial differential equations to be considered in future work.

Acknowledgements We are grateful to the anonymous referees for helpful comments and references.

A Derivative of discrete gradients of the Coordinate Increment type

We here give the expression for the derivative map of the discrete gradients defined in terms of the coordinate increment method of Itoh and Abe [7]. We write (1.3) in compact notation as

$$(\bar{\nabla}_{\text{CI}} H(v, u))_i = \frac{H(u|_i v) - H(u|_{i-1} v)}{u_i - v_i}, \quad 1 \leq i \leq m$$

The jacobian of this map with respect to its second argument is then simply the lower-triangular matrix with elements

$$(D\bar{\nabla}_{\text{CI}} H(v, u))_{ij} = \begin{cases} \frac{\frac{\partial}{\partial u_j}(H(u|_i v) - H(u|_{i-1} v))}{u_i - v_i}, & j < i \\ \frac{\frac{\partial H}{\partial u_i}(u|_i v)}{u_i - v_i} - \frac{H(u|_i v) - H(u|_{i-1} v)}{(u_i - v_i)^2}, & j = i \\ 0 & j > i \end{cases}$$

Similarly, we can compute the Jacobian of the symmetrised version as follows

$$(D\bar{\nabla}_{\text{SCI}} H(v, u))_{ij} = \begin{cases} \frac{1}{2} \frac{\frac{\partial}{\partial u_j}(H(u|_i v) - H(u|_{i-1} v))}{u_i - v_i}, & j < i \\ \frac{1}{2} \frac{\frac{\partial}{\partial u_i}(H(u|_i v) + H(v|_{i-1} u))}{u_i - v_i} + \frac{1}{2} \frac{H(v|_i u) + H(u|_{i-1} v) - H(v|_{i-1} u) - H(u|_i v)}{(u_i - v_i)^2}, & j = i \\ \frac{1}{2} \frac{\frac{\partial}{\partial u_j}(H(v|_{i-1} u) - H(v|_i u))}{u_i - v_i}, & j > i. \end{cases}$$

References

- [1] E. Celledoni and B. Owren. A class of intrinsic schemes for orthogonal integration. *SIAM J. Numer. Anal.*, 40(6):2069–2084 (electronic) (2003), 2002.
- [2] J. L. Cieřliński. An orbit-preserving discretization of the classical Kepler problem. *Phys. Lett. A*, 370(1):8–12, 2007.
- [3] J. L. Cieřliński. Comment on ‘Conservative discretizations of the Kepler motion’. *J. Phys. A*, 43(22):228001, 4, 2010.

- [4] R. Courant, K. Friedrichs, and H. Lewy. Über die partiellen Differenzengleichungen der mathematischen Physik. *Math. Ann.*, 100(1):32–74, 1928.
- [5] O. Gonzalez. Time integration and discrete Hamiltonian systems. *J. Non-linear Sci.*, 6(5):449–467, 1996.
- [6] E. Hairer, C. Lubich, and G. Wanner. *Geometric numerical integration*, volume 31 of *Springer Series in Computational Mathematics*. Springer-Verlag, Berlin, second edition, 2006. Structure-preserving algorithms for ordinary differential equations.
- [7] T. Itoh and K. Abe. Hamiltonian-conserving discrete canonical equations based on variational difference quotients. *J. Comput. Phys.*, 76(1):85–102, 1988.
- [8] R. Kozlov. Conservative discretizations of the Kepler motion. *J. Phys. A*, 40(17):4529–4539, 2007.
- [9] R. A. LaBudde and D. Greenspan. Discrete mechanics—a general treatment. *J. Computational Phys.*, 15:134–167, 1974.
- [10] R. A. LaBudde and D. Greenspan. Energy and momentum conserving methods of arbitrary order of the numerical integration of equations of motion. I. Motion of a single particle. *Numer. Math.*, 25(4):323–346, 1975/76.
- [11] R. A. LaBudde and D. Greenspan. Energy and momentum conserving methods of arbitrary order for the numerical integration of equations of motion. II. Motion of a system of particles. *Numer. Math.*, 26(1):1–16, 1976.
- [12] R. I. McLachlan, G. R. W. Quispel, and N. Robidoux. Geometric integration using discrete gradients. *R. Soc. Lond. Philos. Trans. Ser. A Math. Phys. Eng. Sci.*, 357(1754):1021–1045, 1999.
- [13] Y. Minesaki and Y. Nakamura. A new conservative numerical integration algorithm for the three-dimensional Kepler motion based on the Kustaanheimo-Stiefel regularization theory. *Phys. Lett. A*, 324(4):282–292, 2004.
- [14] Y. Minesaki and Y. Nakamura. New numerical integrator for the Stäckel system conserving the same number of constants of motion as the degree of freedom. *J. Phys. A*, 39(30):9453–9476, 2006.
- [15] F. A. Potra and W. C. Rheinboldt. On the numerical solution of Euler-Lagrange equations. *Mech. Structures Mach.*, 19(1):1–18, 1991.
- [16] F. A. Potra and J. Yen. Implicit numerical integration for Euler-Lagrange equations via tangent space parametrization. *Mech. Structures Mach.*, 19(1):77–98, 1991.
- [17] G. R. W. Quispel and D. I. McLaren. A new class of energy-preserving numerical integration methods. *J. Phys. A*, 41(4):045206, 7, 2008.

- [18] G.R.W. Quispel and H. Capel. Solving ODE's numerically while preserving all first integrals, 1999. Preprint (http://www.latrobe.edu.au/mathstats/staff/quispel/quispel/Publ55_Solving%20ODE's%20numerically.pdf).
- [19] J. C. Simo, N. Tarnow, and K. K. Wong. Exact energy-momentum conserving algorithms and symplectic schemes for nonlinear dynamics. *Comput. Methods Appl. Mech. Engrg.*, 100(1):63–116, 1992.
- [20] E. L. Stiefel and G. Scheifele. *Linear and regular celestial mechanics. Perturbed two-body motion, numerical methods, canonical theory*. Springer-Verlag, New York, 1971. Die Grundlehren der mathematischen Wissenschaften, Band 174.
- [21] L. N. Trefethen and D. Bau. *Numerical linear algebra*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1997.